

Nonasymptotic bounds on the mean square error for MCMC estimates via renewal techniques

Krzysztof Łatuszyński, Błażej Miasojedow and Wojciech Niemiro

Abstract The Nummellin’s split chain construction allows to decompose a Markov chain Monte Carlo (MCMC) trajectory into i.i.d. “excursions”. Regenerative MCMC algorithms based on this technique use a random number of samples. They have been proposed as a promising alternative to usual fixed length simulation [24, 32, 13]. In this note we derive nonasymptotic bounds on the mean square error (MSE) of regenerative MCMC estimates via techniques of renewal theory and sequential statistics. These results are applied to construct confidence intervals. We then focus on two cases of particular interest: chains satisfying the Doeblin condition and a geometric drift condition. Available explicit nonasymptotic results are compared for different schemes of MCMC simulation.

1 Introduction

Consider a typical MCMC setting, where π is a probability distribution on \mathcal{X} and $f : \mathcal{X} \rightarrow \mathbb{R}$ a Borel measurable function. The objective is to compute (estimate) the integral

Krzysztof Łatuszyński
Department of Statistics
University of Warwick
CV4 7AL, Coventry, UK, e-mail: latuch@gmail.com

Błażej Miasojedow
Institute of Applied Mathematics and Mechanics
University of Warsaw
Banacha 2, 02-097 Warszawa, Poland, e-mail: bmia@mimuw.edu.pl

Wojciech Niemiro
Faculty of Mathematics and Computer Science
Nicolaus Copernicus University
Chopina 12/18, 87-100 Toruń, Poland, e-mail: wniemiro@gmail.com

$$\theta := \pi f = \int_{\mathcal{X}} \pi(\mathrm{d}x) f(x). \quad (1)$$

Assume that direct simulation from π is intractable. Therefore one uses an ergodic Markov chain with transition kernel P and stationary distribution π to sample approximately from π . Numerous computational problems from Bayesian inference, statistical physics or combinatorial enumeration fit into this setting. We refer to [31, 29, 9] for theory and applications of MCMC.

Let $(X_n)_{n \geq 0}$ be the Markov chain in question. Typically one discards an initial part of the trajectory (called burn-in, say of length t) to reduce bias, one simulates the chain for n further steps and one approximates θ with an ergodic average:

$$\hat{\theta}_{t,n}^{\text{fix}} = \frac{1}{n} \sum_{i=t}^{t+n-1} f(X_i). \quad (2)$$

The fixed numbers t and n are the parameters of the algorithm. Asymptotic validity of (2) is ensured by a Strong Law of Large Numbers and a Central Limit Theorem (CLT). Under appropriate regularity conditions [31, 4], it holds that

$$\sqrt{n}(\hat{\theta}_{t,n}^{\text{fix}} - \theta) \rightarrow \mathcal{N}(0, \sigma_{\text{as}}^2(f)), \quad (n \rightarrow \infty), \quad (3)$$

where $\sigma_{\text{as}}^2(f)$ is called the asymptotic variance. In contrast with the asymptotic theory, explicit *nonasymptotic* error bounds for $\hat{\theta}_{t,n}^{\text{fix}}$ appear to be very difficult to derive in practically meaningful problems.

Regenerative simulation offers a way to get around some of the difficulties. The split chain construction introduced in [2, 27] (to be described in Section 2) allows for partitioning the trajectory $(X_n)_{n \geq 0}$ into i.i.d. random tours (excursions) between consecutive regeneration times T_0, T_1, T_2, \dots . Random variables

$$\Xi_k(f) := \sum_{i=T_{k-1}}^{T_k-1} f(X_i) \quad (4)$$

are i.i.d. for $k = 1, 2, \dots$ ($\Xi_0(f)$ can have a different distribution). Mykland et al. in [24] suggested a practically relevant recipe for identifying T_0, T_1, T_2, \dots in simulations (formula (2) in Section 2). This resolves the burn-in problem since one can just ignore the part until the first regeneration T_0 . One can also stop the simulation at a regeneration time, say T_r , and simulate r full i.i.d. tours, c.f. Section 4 of [32]. Thus one estimates θ by

$$\hat{\theta}_r^{\text{reg}} := \frac{1}{T_r - T_0} \sum_{i=T_0}^{T_r-1} f(X_i) = \frac{\sum_{k=1}^r \Xi_k(f)}{\sum_{k=1}^r \tau_k}, \quad (5)$$

where $\tau_k = T_k - T_{k-1} = \Xi_k(1)$ are the lengths of excursions. The number of tours r is fixed and the total simulation effort T_r is random. Since $\hat{\theta}_r^{\text{reg}}$ involves i.i.d. random variables, classical tools seem to be sufficient to analyse its behaviour. Asymptotically, (5) is equivalent to (2) because

$$\sqrt{rm}(\hat{\theta}_r^{\text{reg}} - \theta) \rightarrow \mathcal{N}(0, \sigma_{\text{as}}^2(f)), \quad (r \rightarrow \infty),$$

where $m := \mathbb{E} \tau_1$. Now $rm = \mathbb{E}(T_r - T_0)$, the expected length of the trajectory, plays the role of n . However, our attempt at nonasymptotic analysis in Subsection 3.1 reveals unexpected difficulties: our bounds involve m in the denominator and in most practically relevant situations m is unknown.

If m is known then instead of (5) one can use an unbiased estimator

$$\tilde{\theta}_r^{\text{unb}} := \frac{1}{rm} \sum_{k=1}^r \Xi_k(f), \quad (6)$$

Quite unexpectedly, (6) is *not equivalent* to (5), even in a weak asymptotic sense. The standard CLT for i.i.d. summands yields

$$\sqrt{rm}(\tilde{\theta}_r^{\text{unb}} - \theta) \rightarrow \mathcal{N}(0, \sigma_{\text{unb}}^2(f)), \quad (r \rightarrow \infty),$$

where $\sigma_{\text{unb}}^2(f) := \text{Var} \Xi_1(f)/m$ is in general different from $\sigma_{\text{as}}^2(f)$.

We introduce a new regenerative-sequential simulation scheme, for which better nonasymptotic results can be derived. Namely, we fix n and define

$$R(n) := \min\{r : T_r > T_0 + n\}.$$

The estimator is defined as

$$\hat{\theta}_n^{\text{reg-seq}} := \frac{1}{T_{R(n)} - T_0} \sum_{i=T_0}^{T_{R(n)}-1} f(X_i) = \frac{\sum_{k=1}^{R(n)} \Xi_k(f)}{\sum_{k=1}^{R(n)} \tau_k}. \quad (7)$$

We thus generate a random number of tours as well as a random number of samples.

Our approach is based on inequalities for the mean square error,

$$\text{MSE} := \mathbb{E}(\hat{\theta} - \theta)^2.$$

Bounds on the MSE can be used to construct fixed precision confidence intervals. The goal is to obtain an estimator $\hat{\theta}$ which satisfies

$$\mathbb{P}(|\hat{\theta} - \theta| \leq \varepsilon) \geq 1 - \alpha, \quad (8)$$

for given ε and α . We combine the MSE bounds with the so called “median trick” [15, 26]. One runs MCMC repeatedly and computes the median of independent estimates to boost the level of confidence. In our paper, the median trick is used in conjunction with regenerative simulation.

The organization of the paper is the following. In Section 2 we recall the split chain construction. Nonasymptotic bounds for regenerative estimators defined by (5), (6) and (7) are derived in Section 3. Derivation of more explicit bounds which involve only computable quantities is deferred to Sections 5 and 6, where we consider classes of chains particularly important in the MCMC context. An analogous

analysis of the non-regenerative scheme (2) was considered in [20] and (in a different setting and using different methods) in [33].

In Section 4 we discuss the median trick. The resulting confidence intervals are compared with *asymptotic* results based on the CLT.

In Section 5 we consider Doeblin chains, i.e. uniformly ergodic chains that satisfy a one step minorization condition. We compare regenerative estimators (5), (6) and (7). Moreover, we also consider a perfect sampler available for Doeblin chains, c.f. [35, 14]. We show that confidence intervals based on the median trick can outperform those obtained via exponential inequalities for a single run simulation.

In Section 6 we proceed to analyze geometrically ergodic Markov chains, assuming a drift condition towards a small set. We briefly compare regenerative schemes (5) and (7) in this setting (the unbiased estimator (6) cannot be used, because m is unknown).

2 Regenerative Simulation

We describe the setting more precisely. Let $(X_n)_{n \geq 0}$ be a Markov chain with transition kernel P on a Polish space \mathcal{X} with stationary distribution π , i.e. $\pi P = \pi$. Assume P is π -irreducible. The regeneration/split construction of Nummelin [27] and Athreya and Ney [2] rests on the following assumption.

Assumption 2.1 (Small Set) *There exist a Borel set $J \subseteq \mathcal{X}$ of positive π measure, a number $\beta > 0$ and a probability measure ν such that*

$$P(x, \cdot) \geq \beta \mathbb{I}(x \in J) \nu(\cdot).$$

Under Assumption 2.1 we can define a bivariate Markov chain (X_n, Γ_n) on the space $\mathcal{X} \times \{0, 1\}$ in the following way. Variable Γ_{n-1} depends only on X_{n-1} via $\mathbb{P}(\Gamma_{n-1} = 1 | X_{n-1} = x) = \beta \mathbb{I}(x \in J)$. The rule of transition from (X_{n-1}, Γ_{n-1}) to X_n is given by

$$\begin{aligned} \mathbb{P}(X_n \in A | \Gamma_{n-1} = 1, X_{n-1} = x) &= \nu(A), \\ \mathbb{P}(X_n \in A | \Gamma_{n-1} = 0, X_{n-1} = x) &= Q(x, A), \end{aligned}$$

where Q is the normalized “residual” kernel given by

$$Q(x, \cdot) := \frac{P(x, \cdot) - \beta \mathbb{I}(x \in J) \nu(\cdot)}{1 - \beta \mathbb{I}(x \in J)}.$$

Whenever $\Gamma_{n-1} = 1$, the chain regenerates at moment n . The regeneration epochs are

$$\begin{aligned} T_0 &:= \min\{n : \Gamma_{n-1} = 1\}, \\ T_k &:= \min\{n > T_{k-1} : \Gamma_{n-1} = 1\}. \end{aligned}$$

The random tours defined by

$$\Xi_k := (X_{T_{k-1}}, \dots, X_{T_k-1}, \tau_k), \quad \text{where } \tau_k = T_k - T_{k-1}, \quad (9)$$

are independent. Without loss of generality, we assume that $X_0 \sim \nu(\cdot)$, unless stated otherwise. Under this assumption, all the tours Ξ_k are i.i.d. for $k > 0$. We therefore put $T_0 := 0$ and simplify notation. In the sequel symbols \mathbb{P} and \mathbb{E} without subscripts refer to the chain started at ν . If the initial distribution ξ is other than ν , it will be explicitly indicated by writing \mathbb{P}_ξ and \mathbb{E}_ξ . Notation $m = \mathbb{E} \tau_1$ stands throughout the paper.

We assume that we are able to *identify* regeneration times T_k . Mykland et al. pointed out in [24] that actual sampling from Q can be avoided. We can generate the chain using transition probability P and then recover the regeneration indicators via

$$\mathbb{P}(\Gamma_{n-1} = 1 | X_n, X_{n-1}) = \mathbb{I}(X_{n-1} \in J) \frac{\beta \nu(dX_n)}{P(X_{n-1}, dX_n)},$$

where $\nu(dy)/P(x, dy)$ denote the Radon-Nikodym derivative (in practice, the ratio of densities). Mykland's trick has been established in a number of practically relevant families (e.g. hierarchical linear models) and specific Markov chains implementations, such as block Gibbs samplers or variable-at-a-time chains, see [17, 25].

3 General results for regenerative estimators

Recall that $f : \mathcal{X} \rightarrow \mathbb{R}$ is a measurable function and $\theta = \pi f$. We consider block sums $\Xi_k(f)$ defined by (4). The general Kac theorem states that the mean occupation time during one tour is proportional to the stationary measure (Theorem 10.0.1 in [23] or Equations (3.3.4), (3.3.6), (3.4.7) and (3.5.1) in [28]). This yields

$$m = \frac{1}{\beta \pi(J)}, \quad \mathbb{E} \Xi_1(f) = m \pi f = m \theta.$$

From now on we assume that $\mathbb{E} \Xi_1(f)^2 < \infty$ and $\mathbb{E} \tau_1^2 < \infty$. For a discussion of these assumptions in the MCMC context, see [13]. Let $\bar{f} := f - \pi f$ and define

$$\sigma_{\text{as}}^2(f) := \frac{\mathbb{E} \Xi_1(\bar{f})^2}{m}, \quad (10)$$

$$\sigma_\tau^2 := \frac{\text{Var} \tau_1}{m}. \quad (11)$$

Remark 1. Under Assumption 2.1, finiteness of $\mathbb{E} \Xi_1(\bar{f})^2$ is a sufficient and necessary condition for the CLT to hold for Markov chain $(X_n)_{n \geq 0}$ and function f . This fact is proved in [4] in a more general setting. For our purposes it is important to note that $\sigma_{\text{as}}^2(f)$ in (10) is indeed the *asymptotic variance* which appears in the CLT.

3.1 Results for $\hat{\theta}_r^{\text{reg}}$

We are to bound the estimation error which can be expressed as follows:

$$\hat{\theta}_r^{\text{reg}} - \theta = \frac{\sum_{k=1}^r (\Xi_k(f) - \theta \tau_k)}{\sum_{k=1}^r \tau_k} = \frac{\sum_{k=1}^r d_k}{T_r}. \quad (12)$$

where $d_k := \Xi_k(f) - \theta \tau_k = \Xi_k(\bar{f})$. Therefore, for any $0 < \delta < 1$,

$$\mathbb{P}(|\hat{\theta}_r^{\text{reg}} - \theta| > \varepsilon) \leq \mathbb{P}\left(\left|\sum_{k=1}^r d_k\right| > rm\varepsilon(1 - \delta)\right) + \mathbb{P}(T_r < rm(1 - \delta)).$$

Since d_k are i.i.d. with $\mathbb{E}d_1 = 0$ and $\text{Var}d_1 = m\sigma_{\text{as}}^2(f)$, we can use Chebyshev inequality to bound the first term above:

$$\mathbb{P}\left(\left|\sum_{k=1}^r d_k\right| > rm\varepsilon(1 - \delta)\right) \leq \frac{\sigma_{\text{as}}^2(f)}{rm\varepsilon^2(1 - \delta)^2}.$$

The second term can be bounded similarly. We use the fact that τ_k are i.i.d. with $\mathbb{E}\tau_1 = m$ to write

$$\mathbb{P}(T_r < rm(1 - \delta)) \leq \frac{\sigma_{\tau}^2}{rm^2\delta^2}.$$

We conclude the above calculation with in following Theorem.

Theorem 3.1 *Under Assumption 2.1 the following holds for every $0 < \delta < 1$*

$$\mathbb{P}(|\hat{\theta}_r^{\text{reg}} - \theta| > \varepsilon) \leq \frac{1}{rm} \left[\frac{\sigma_{\text{as}}^2(f)}{\varepsilon^2(1 - \delta)^2} + \frac{\sigma_{\tau}^2}{m\delta^2} \right] \quad (13)$$

and is minimized by

$$\delta = \delta^* := \frac{\sigma_{\tau}^{2/3}}{\sigma_{\text{as}}^{2/3}(f)\varepsilon^{-2/3} + \sigma_{\tau}^{2/3}}.$$

Obviously, $\mathbb{E}T_r = rm$ is the expected length of trajectory. The main drawback of Theorem 3.1 is that the bound on the estimation error depends on m , which is typically unknown. Replacing m by 1 in (13) would be highly inefficient. This fact motivates our study of another estimator, $\hat{\theta}_n^{\text{reg-seq}}$, for which we can obtain much more satisfactory results. We think that the derivation of better nonasymptotic bounds for $\hat{\theta}_r^{\text{reg}}$ (not involving m) is an open problem.

3.2 Results for $\tilde{\theta}_r^{\text{unb}}$

Recall that $\tilde{\theta}_r^{\text{unb}}$ can be used only when m is known and this situation is rather rare in MCMC applications. The analysis of $\tilde{\theta}_r^{\text{unb}}$ is straightforward, because it is simply a sum of i.i.d. random variables. In particular, we obtain the following.

Corollary 3.2 *Under Assumption 2.1,*

$$\mathbb{E}(\tilde{\theta}_r^{\text{unb}} - \theta)^2 = \frac{\sigma_{\text{unb}}^2(f)}{rm}, \quad \mathbb{P}(|\tilde{\theta}_r^{\text{unb}} - \theta| > \varepsilon) \leq \frac{\sigma_{\text{unb}}^2(f)}{rm\varepsilon^2}.$$

Note that $\sigma_{\text{unb}}^2(f) = \text{Var}\Xi_1(f)/m$ can be expressed as

$$\sigma_{\text{unb}}^2(f) = \sigma_{\text{as}}^2(f) + \theta^2 \sigma_{\tau}^2 + 2\theta\rho(\bar{f}, 1), \quad (14)$$

where $\rho(\bar{f}, 1) := \text{Cov}(\Xi_1(\bar{f}), \Xi_1(1))/m$. This follows from the simple observation that $\text{Var}\Xi_1(f) = \mathbb{E}(\Xi_1(\bar{f}) + \theta(\tau_1 - m))^2$.

3.3 Results for $\hat{\theta}_n^{\text{reg-seq}}$

The result below bounds the MSE and the expected number of samples used to compute the estimator.

Theorem 3.3 *If Assumption 2.1 holds then*

$$(i) \quad \mathbb{E}(\hat{\theta}_n^{\text{reg-seq}} - \theta)^2 \leq \frac{\sigma_{\text{as}}^2(f)}{n^2} \mathbb{E} T_{R(n)}$$

and

$$(ii) \quad \mathbb{E} T_{R(n)} \leq n + C_0,$$

where

$$C_0 := \sigma_{\tau}^2 + m.$$

Corollary 3.4 *Under Assumption 2.1,*

$$\mathbb{E}(\hat{\theta}_n^{\text{reg-seq}} - \theta)^2 \leq \frac{\sigma_{\text{as}}^2(f)}{n} \left(1 + \frac{C_0}{n}\right), \quad (15)$$

$$\mathbb{P}(|\hat{\theta}_n^{\text{reg-seq}} - \theta| > \varepsilon) \leq \frac{\sigma_{\text{as}}^2(f)}{n\varepsilon^2} \left(1 + \frac{C_0}{n}\right). \quad (16)$$

Remark 2. Note that the leading term $\sigma_{\text{as}}^2(f)/n$ in (15) is “asymptotically correct” in the sense that the standard fixed length estimator has $\text{MSE} \sim \sigma_{\text{as}}^2(f)/n$. The regenerative-sequential scheme is “close to the fixed length simulation”, because $\lim_{n \rightarrow \infty} \mathbb{E} T_{R(n)}/n = 1$.

Proof (of Theorem 3.3). Just as in (12) we have

$$\hat{\theta}_n^{\text{reg-seq}} - \theta = \frac{\sum_{k=1}^{R(n)} (\Xi_k(f) - \theta \tau_k)}{\sum_{k=1}^{R(n)} \tau_k} = \frac{1}{T_{R(n)}} \sum_{k=1}^{R(n)} d_k,$$

where pairs (d_k, τ_k) are i.i.d. with $\mathbb{E} d_1 = 0$ and $\text{Var} d_1 = m\sigma_{\text{as}}^2(f)$. Since $T_{R(n)} > n$, it follows that

$$\mathbb{E} (\hat{\theta}_n^{\text{reg-seq}} - \theta)^2 \leq \frac{1}{n^2} \mathbb{E} \left(\sum_{k=1}^{R(n)} d_k \right)^2.$$

Since $R(n)$ is a stopping time with respect to $\mathcal{G}_k = \sigma((d_1, \tau_1), \dots, (d_k, \tau_k))$, we are in a position to apply the two Wald's identities (see Appendix). The second identity yields

$$\mathbb{E} \left(\sum_{k=1}^{R(n)} d_k \right)^2 = \text{Var} d_1 \mathbb{E} R(n) = m\sigma_{\text{as}}^2(f) \mathbb{E} R(n).$$

In this expression we can replace $m\mathbb{E} R(n)$ by $\mathbb{E} T_{R(n)}$ because of the first Wald's identity:

$$\mathbb{E} T_{R(n)} = \mathbb{E} \sum_{k=1}^{R(n)} \tau_k = \mathbb{E} \tau_1 \mathbb{E} R(n) = m\mathbb{E} R(n)$$

and (i) follows.

We now focus attention on bounding the expectation of the “overshoot” $\Delta(n) := T_{R(n)} - n$. Since we assume that $X_0 \sim \nu$, the cumulative sums $\tau_1 = T_1 < T_2 < \dots < T_k < \dots$ form a (nondelayed) renewal process in discrete time. Let us invoke the following elegant theorem of Lorden [21, Theorem 1]:

$$\mathbb{E} \Delta(n) \leq \mathbb{E} \tau_1^2 / m.$$

This inequality combined with (11) yields immediately $\mathbb{E} T_{R(n)} = \mathbb{E} (n + \Delta(n)) \leq n + \sigma_{\tau}^2 + m$, i.e. (ii).

4 The median trick

This ingenious method of constructing fixed precision MCMC algorithms was introduced in 1986 in [15], later used in many papers concerned with computational complexity and further developed in [26]. We run l independent copies of the Markov chain. Let $\hat{\theta}^{(j)}$ be an estimator computed in j th run. The final estimate is $\hat{\theta} := \text{med}(\hat{\theta}^{(1)}, \dots, \hat{\theta}^{(l)})$. To ensure that $\hat{\theta}$ satisfies (8), we require that $\mathbb{P}(|\hat{\theta}^{(j)} - \theta| > \varepsilon) \leq a$ ($j = 1, \dots, l$) for some modest level of confidence $1 - a < 1 - \alpha$. This is obtained via Chebyshev's inequality, if a bound on MSE is available. The well-known Chernoff's inequality gives for odd l ,

$$\mathbb{P}(|\hat{\theta} - \theta| \geq \varepsilon) \leq \frac{1}{2} [4a(1-a)]^{l/2} = \frac{1}{2} \exp \left\{ \frac{l}{2} \ln [4a(1-a)] \right\}. \quad (17)$$

It is pointed out in [26] that under some assumptions there is a universal choice of a , which nearly minimizes the overall number of samples, $a^* \approx 0.11969$.

Let us now examine how the median trick works in conjunction with regenerative MCMC. We focus on $\hat{\theta}_n^{\text{reg-seq}}$, because Corollary 3.4 gives the best available bound on MSE. We first choose n such that the right hand side of (16) is less than or equal to a^* . Then choose l big enough to make the right hand side of (17) (with $a = a^*$) less than or equal to α . Compute estimator $\hat{\theta}_n^{\text{reg-seq}}$ repeatedly, using l independent runs of the chain. We can see that (8) holds if

$$n \geq \frac{C_1 \sigma_{\text{as}}^2(f)}{\varepsilon^2} + C_0, \quad (18)$$

$$l \geq C_2 \ln(2\alpha)^{-1} \text{ and } l \text{ is odd}, \quad (19)$$

where $C_1 := 1/a^* \approx 8.3549$ and $C_2 := 2/\ln[4a^*(1-a^*)]^{-1} \approx 2.3147$ are absolute constants. Indeed, (18) entails $C_1 \sigma_{\text{as}}^2(f)/(\varepsilon^2 n) \leq 1 - C_0/n$, so $C_1 \sigma_{\text{as}}^2(f)/(\varepsilon^2 n)(1 + C_0/n) \leq 1 - C_0^2/n^2 < 1$. Consequently $\sigma_{\text{as}}^2(f)/(\varepsilon^2 n)(1 + C_0/n) < a^*$ and we are in a position to apply (16).

The overall (expected) number of generated samples is $l\mathbb{E}T_{R(n)} \sim nl$ as $\varepsilon \rightarrow 0$ and $n \rightarrow \infty$, by Theorem 3.3 (ii). Consequently for $\varepsilon \rightarrow 0$ the cost of the algorithm is approximately

$$nl \sim C \frac{\sigma_{\text{as}}^2(f)}{\varepsilon^2} \log(2\alpha)^{-1}, \quad (20)$$

where $C = C_1 C_2 \approx 19.34$. To see how tight is the obtained lower bound, let us compare (20) with the familiar asymptotic approximation, based on the CLT. Consider an estimator based on one MCMC run of length n , say $\hat{\theta}_n = \hat{\theta}_{t,n}^{\text{fix}}$ with $t = 0$. From (3) we infer that

$$\lim_{\varepsilon \rightarrow 0} \mathbb{P}(|\hat{\theta}_n - \theta| > \varepsilon) = \alpha,$$

holds for

$$n \sim \frac{\sigma_{\text{as}}^2(f)}{\varepsilon^2} [\Phi^{-1}(1 - \alpha/2)]^2, \quad (21)$$

where Φ^{-1} is the quantile function of the standard normal distribution. Taking into account the fact that $[\Phi^{-1}(1 - \alpha/2)]^2 \sim 2\log(2\alpha)^{-1}$ for $\alpha \rightarrow 0$ we arrive at the following conclusion. The right hand side of (20) is bigger than (21) roughly by a constant factor of about 10 (for small ε and α). The important difference is that (20) is sufficient for an *exact* confidence interval while (21) only for an *asymptotic* one.

5 Doeblin Chains

Assume that the transition kernel P satisfies the following Doeblin condition: there exist $\beta > 0$ and a probability measure ν such that

$$P(x, \cdot) \geq \beta \nu(\cdot) \quad \text{for every } x \in \mathcal{X}. \quad (22)$$

This amounts to taking $J := \mathcal{X}$ in Assumption 2.1. Condition (22) implies that the chain is uniformly ergodic. We refer to [31] and [23] for definition of uniform ergodicity and related concepts. As a consequence of the regeneration construction, in our present setting τ_1 is distributed as a geometric random variable with parameter β and therefore

$$m = \mathbb{E} \tau_1 = \frac{1}{\beta} \quad \text{and} \quad \sigma_\tau^2 = \frac{\text{Var} \tau_1}{m} = \frac{1 - \beta}{\beta}.$$

Bounds on the asymptotic variance $\sigma_{\text{as}}^2(f)$ under (22) are well known. Let $\sigma^2 = \pi \bar{f}^2$ be the stationary variance. Results in Section 5 of [4] imply that

$$\sigma_{\text{as}}^2(f) \leq \sigma^2 \left(1 + \frac{2\sqrt{1-\beta}}{1-\sqrt{1-\beta}} \right) \leq \frac{4\sigma^2}{\beta}. \quad (23)$$

Since in [4] a more general situation is considered, which complicates the formulas, let us give a simple derivation of (23) under (22). By (10) and the formula (29) given in the Appendix,

$$\sigma_{\text{as}}^2(f) \leq \frac{\mathbb{E} \Xi_1(|\bar{f}|)^2}{m} = \mathbb{E} \pi \bar{f}(X_0)^2 + 2 \sum_{i=1}^{\infty} \mathbb{E} \pi |\bar{f}(X_0) \bar{f}(X_i)| \mathbb{I}(\tau_1 > i).$$

The first term above is equal to σ^2 . To bound the terms of the series, use Cauchy-Schwarz and the fact that, under (22), random variables X_0 and τ_1 are independent. Therefore $\mathbb{E} \pi |\bar{f}(X_0) \bar{f}(X_i)| \mathbb{I}(\tau_1 > i) \leq (\mathbb{E} \pi \bar{f}(X_i)^2 \mathbb{E} \pi \bar{f}(X_0)^2 \mathbb{P}_\pi(\tau_1 > i))^{1/2} = \sigma^2 (1 - \beta)^{i/2}$. Computing the sum of the geometric series yields (23).

If the chain is reversible, there is a better bound than (23). We can use the well-known formula for $\sigma_{\text{as}}^2(f)$ in terms of the spectral decomposition of P (e.g. expression “C” in [11]). Results of [30] show that the spectrum of P is a subset of $[-1 + \beta, 1 - \beta]$. We conclude that for reversible Doeblin chains,

$$\sigma_{\text{as}}^2(f) \leq \frac{2 - \beta}{\beta} \sigma^2 \leq \frac{2\sigma^2}{\beta}. \quad (24)$$

An important class of reversible chains are Independence Metropolis-Hastings chains (see e.g. [31]) that are known to be uniformly ergodic if and only if the rejection probability $r(x)$ is uniformly bounded from 1 by say $1 - \beta$. This is equivalent to the candidate distribution being bounded below by $\beta \pi$ (c.f. [22, 1]) and translates

into (22) with $\nu = \pi$. The formula for $\sigma_{\text{as}}^2(f)$ in (23) and (24) depends on β in an optimal way. Moreover (24) is sharp. To see this consider the following example.

Example 5.1 Let $\beta \leq 1/2$ and define a Markov chain $(X_n)_{n \geq 0}$ on $\mathcal{X} = \{0, 1\}$ with stationary distribution $\pi = [1/2, 1/2]$ and transition matrix

$$P = \begin{bmatrix} 1 - \beta/2 & \beta/2 \\ \beta/2 & 1 - \beta/2 \end{bmatrix}.$$

Hence $P = \beta\pi + (1 - \beta)I_2$ and $P(x, \cdot) \geq \beta\pi$. Note that the residual kernel Q is in our example the identity matrix I_2 . Thus, before the first regeneration τ_1 the chain does not move. Let $f(x) = x$. Thus $\sigma^2 = 1/4$. To compute $\sigma_{\text{as}}^2(f)$ we use another well-known formula (expression “B” in [11]):

$$\begin{aligned} \sigma_{\text{as}}^2(f) &= \sigma^2 + 2 \sum_{i=1}^{\infty} \text{Cov}_{\pi}\{f(X_0), f(X_i)\} \\ &= \sigma^2 + 2\sigma^2 \sum_{i=1}^{\infty} (1 - \beta)^i = \frac{2 - \beta}{\beta} \sigma^2. \end{aligned}$$

To compute $\sigma_{\text{unb}}^2(f)$, note that $\Xi_1(f) = \mathbb{I}(X_0 = 1)\tau_1$. Since τ_1 is independent of X_0 and $X_0 \sim \nu = \pi$ we obtain

$$\begin{aligned} \sigma_{\text{unb}}^2(f) &= \beta \text{Var} \Xi_1(f) = \beta [\mathbb{E} \text{Var}(\Xi_1(f)|X_0) + \text{Var} \mathbb{E}(\Xi_1(f)|X_0)] \\ &= \frac{1 - \beta}{2\beta} + \frac{1}{4\beta} = \frac{3 - 2\beta}{\beta} \sigma^2. \end{aligned}$$

Interestingly, in this example $\sigma_{\text{unb}}^2(f) > \sigma_{\text{as}}^2(f)$.

In the setting of this Section, we will now compare upper bounds on the total simulation effort needed for different MCMC schemes to get $\mathbb{P}(|\hat{\theta} - \theta| > \varepsilon) \leq \alpha$.

5.1 Regenerative-sequential estimator and the median trick

Recall that this simulation scheme consists of l MCMC runs, each of approximate length n . Substituting either (23) or (24) in (20) we obtain that the expected number of samples is

$$nl \sim 19.34 \frac{4\sigma^2}{\beta\varepsilon^2} \log(2\alpha)^{-1} \quad \text{and} \quad nl \sim 19.34 \frac{(2 - \beta)\sigma^2}{\beta\varepsilon^2} \log(2\alpha)^{-1} \quad (25)$$

(respectively in the general case and for reversible chains). Note also that in the setting of this Section we have an exact expression for the constant C_0 in Theorem 3.3. Indeed, $C_0 = 2/\beta - 1$.

5.2 Standard one-run average and exponential inequality

For uniformly ergodic chains a direct comparison of our approach to exponential inequalities [10, 18] is possible. We focus on the result proved in [18] for chains on a countable state space. This inequality is tight in the sense that it reduces to the Hoeffding bound when specialised to the i.i.d. case. For f bounded let $\|f\|_\infty := \sup_{x \in \mathcal{X}} |f(x)|$. Consider the simple average over n Markov chain samples, say $\hat{\theta}_n = \hat{\theta}_{t,n}^{\text{fix}}$ with $t = 0$. For an arbitrary initial distribution ξ we have

$$\mathbb{P}_\xi(|\hat{\theta}_n - \theta| > \varepsilon) \leq 2 \exp \left\{ -\frac{n-1}{2} \left(\frac{2\beta}{\|f\|_\infty} \varepsilon - \frac{3}{n-1} \right)^2 \right\}.$$

After identifying the leading terms we can see that to make the right hand side less than α we need

$$n \sim \frac{\|f\|_\infty^2}{2\beta^2\varepsilon^2} \log(\alpha/2)^{-1} \geq \frac{2\sigma^2}{\beta^2\varepsilon^2} \log(\alpha/2)^{-1}. \quad (26)$$

Comparing (25) with (26) yields a ratio of roughly 40β or 20β respectively. This in particular indicates that the dependence on β in [10, 18] probably can be improved. We note that in examples of practical interest β usually decays exponentially with the dimension of \mathcal{X} and using the regenerative-sequential-median scheme will often result in a lower total simulation cost. Moreover, this approach is valid for an unbounded target function f , in contrast with classical exponential inequalities.

5.3 Perfect sampler and the median trick

For Doeblin chains, if regeneration times can be identified, perfect sampling can be performed easily as a version of read-once algorithm [35]. This is due to the following observation. If condition (22) holds and $X_0 \sim \nu$ then

$$X_{T_k-1}, \quad k = 1, 2, \dots$$

are i.i.d. random variables from π (see [5, 28, 14, 4] for versions of this result). Therefore from each random tour between regeneration times one can obtain a single perfect sample (by taking the state of the chain prior to regeneration) and use it for i.i.d. estimation. We define

$$\hat{\theta}_r^{\text{perf}} := \frac{1}{r} \sum_{k=1}^r f(X_{T_k-1}).$$

Clearly

$$\mathbb{E}(\hat{\theta}_r^{\text{perf}} - \theta)^2 = \frac{\sigma^2}{r} \quad \text{and} \quad \mathbb{P}(|\hat{\theta}_r^{\text{perf}} - \theta| > \varepsilon) \leq \frac{\sigma^2}{r\varepsilon^2}.$$

Note that to compute $\hat{\theta}_r^{\text{perf}}$ we need to simulate $n \sim r/\beta$ steps of the Markov chain. If we combine the perfect sampler with the median trick we obtain an algorithm with the expected number of samples

$$nl \sim 19.34 \frac{\sigma^2}{\beta \varepsilon^2} \log(2\alpha)^{-1}. \quad (27)$$

Comparing (25) with (26) and (27) leads to the conclusion that if one targets rigorous nonasymptotic results in the Doeblin chain setting, the approach described here outperforms other methods.

5.4 Remarks on other schemes

The bound for $\hat{\theta}_r^{\text{reg}}$ in Theorem 3.1 is clearly inferior to that for $\hat{\theta}_n^{\text{reg-seq}}$ in Corollary 3.4. Therefore we excluded the scheme based on $\hat{\theta}_r^{\text{reg}}$ from our comparisons.

As for $\hat{\theta}_r^{\text{unb}}$, this estimator can be used in the Doeblin chains setting, because $m = 1/\beta$ is known. The bounds for $\hat{\theta}_r^{\text{unb}}$ in Subsection 3.2 involve $\sigma_{\text{unb}}^2(f)$. Although we cannot provide a rigorous proof, we conjecture that in most practical situations we have $\sigma_{\text{unb}}^2(f) > \sigma_{\text{as}}^2(f)$, because $\rho(\bar{f}, 1)$ in (14) is often close to zero. If this is the case, then the bound for $\hat{\theta}_r^{\text{unb}}$ is inferior to that for $\hat{\theta}_n^{\text{reg-seq}}$.

6 A Geometric Drift Condition

Using drift conditions is a standard approach for establishing geometric ergodicity. We refer to [31] or [23] for the definition and further details. The assumption below is the same as in [3]. Specifically, let J be the small set which appears in Assumption 2.1.

Assumption 6.1 (Drift) *There exist a function $V : \mathcal{X} \rightarrow [1, \infty[$, constants $\lambda < 1$ and $K < \infty$ such that*

$$PV(x) := \int_{\mathcal{X}} P(x, dy) V(y) \leq \begin{cases} \lambda V(x) & \text{for } x \notin J, \\ K & \text{for } x \in J, \end{cases}$$

In many papers conditions similar to Assumption 6.1 have been established for realistic MCMC algorithms in statistical models of practical relevance [12, 7, 8, 16, 17, 34]. This opens the possibility of computing our bounds in these models.

Under Assumption 6.1, it is possible to bound $\sigma_{\text{as}}^2(f)$, σ_τ^2 and C_0 which appear in Theorems 3.1 and 3.3, by expressions involving only λ , β and K . The following result is a minor variation of Theorem 6.5 in [19].

Theorem 6.2 *If Assumptions 2.1 and 6.1 hold and f is such that $\|\tilde{f}\|_{V^{1/2}} := \sup_x |\tilde{f}(x)|/V^{1/2}(x) < \infty$, then*

$$\sigma_{\text{as}}^2(f) \leq \|\tilde{f}\|_{V^{1/2}}^2 \left[\frac{1 + \lambda^{1/2}}{1 - \lambda^{1/2}} \pi(V) + \frac{2(K^{1/2} - \lambda^{1/2} - \beta(2 - \lambda^{1/2}))}{\beta(1 - \lambda^{1/2})} \pi(V^{1/2}) \right]$$

$$C_0 \leq \frac{\lambda^{1/2}}{1 - \lambda^{1/2}} \pi(V^{1/2}) + \frac{K^{1/2} - \lambda^{1/2} - \beta}{\beta(1 - \lambda^{1/2})} - 1.$$

To bound σ_τ^2 we can use the obvious inequality $\sigma_\tau^2 = C_0 - m \leq C_0 - 1$. Moreover, one can easily control πV and $\pi V^{1/2}$ and further replace $\|\tilde{f}\|_{V^{1/2}}$ e.g. by $\|f\|_{V^{1/2}} + (K^{1/2} - \lambda^{1/2})/(1 - \lambda^{1/2})$, we refer to [19] for details.

Let us now discuss possible approaches to confidence estimation in the setting of this section. Perfect sampling is in general unavailable. For unbounded f we cannot apply exponential inequalities for the standard one-run estimate. Since m is unknown we cannot use $\hat{\theta}_r^{\text{unb}}$. This leaves $\hat{\theta}_r^{\text{reg}}$ and $\hat{\theta}_n^{\text{reg-seq}}$ combined with the median trick. To analyse $\hat{\theta}_r^{\text{reg}}$ we can apply Theorem 3.1. Upper bounds for $\sigma_{\text{as}}^2(f)$ and σ_τ^2 are available. However, in Theorem 3.1 we will also need a *lower bound* on m . Without further assumptions we can only write

$$m = \frac{1}{\pi(J)\beta} \geq \frac{1}{\beta}. \quad (28)$$

In the above analysis (28) is particularly disappointing. It multiplies the bound by an unexpected and substantial factor, as $\pi(J)$ is typically small in applications. For $\hat{\theta}_n^{\text{reg-seq}}$ we have much more satisfactory results. Theorems 3.3 and 6.2 can be used to obtain bounds which do not involve m . In many realistic examples, the parameters β , λ and K which appear in Assumptions 2.1 (Small Set) and 6.1 (Drift) can be explicitly computed, see e.g. [16, 17, 34].

We note that nonasymptotic confidence intervals for MCMC estimators under drift condition have also been obtained in [20], where identification of regeneration times has not been assumed. In absence of regeneration times a different approach has been used and the bounds are typically weaker. For example one can compare [20, Corollary 3.2] (for estimator $\hat{\theta}_{t,n}^{\text{fix}}$) combined with the bounds in [3] with our Theorems 3.3 and 6.2 (for estimator $\hat{\theta}_n^{\text{reg-seq}}$).

References

1. Y.F. Atchade, F. Perron (2007): On the geometric ergodicity of Metropolis-Hastings algorithms. *Statistics* 41, 77–84.
2. K.B. Athreya and P. Ney (1978): A new approach to the limit theory of recurrent Markov chains, *Trans. Amer. Math. Soc.* 245, 493–501.
3. P.H. Baxendale (2005): Renewal Theory and Computable Convergence Rates for Geometrically Ergodic Markov Chains. *Ann. Appl. Probab.* 15, 700–738.
4. W. Bednorz, R. Latała and K. Łatuszyński (2008): A Regeneration Proof of the Central Limit Theorem for Uniformly Ergodic Markov Chains. *Elect. Comm. in Probab.* 13, 85–98.
5. L.A. Breyer and G.O. Roberts (2001): Catalytic perfect simulation. *Methodol. Comput. Appl. Probab.* 3 161–177.
6. Y.S. Chow and H. Teicher (1988): *Probability Theory, Independence, Interchangeability, Martingales*. Second Edition, Springer Verlag.
7. G. Fort and E. Moulines (2000): V-subgeometric ergodicity for a Hastings–Metropolis algorithm. *Statist. Probab. Lett.* 49, 401–410.
8. G. Fort, E. Moulines, G.O. Roberts, and J.S. Rosenthal (2003): On the geometric ergodicity of hybrid samplers. *J. Appl. Probab.* 40 (1), 123–146.
9. W.R. Gilks, S. Richardson, D.J. Spiegelhalter: *Markov chain Monte Carlo in practice*. Chapman & Hall, 1998.
10. P.W. Glynn and D. Ormoneit (2002): Hoeffding’s inequality for uniformly ergodic Markov chains, *Statist. Probab. Lett.* 56, 143–146.
11. O. Häggström J.S. Rosenthal (2007): On variance conditions for Markov chain CLTs. *Elect. Comm. in Probab.* 12 , 454–464.
12. J.P. Hobert and C.J. Geyer (1998): Geometric ergodicity of Gibbs and block Gibbs samplers for Hierarchical Random Effects Model. *J. Multivariate Anal.* 67, 414–439.
13. J.P. Hobert, G.L. Jones, B. Presnell, and J.S. Rosenthal (2002): On the Applicability of Regenerative Simulation in Markov Chain Monte Carlo. *Biometrika* 89, 731–743.
14. J.P. Hobert and C.P. Robert (2004): A mixture representation of π with applications in Markov chain Monte Carlo and perfect sampling. *Ann. Appl. Probab.* 14 1295–1305.
15. M.R. Jerrum, L.G. Valiant, V.V. Vazirani (1986): Random generation of combinatorial structures fro, a uniform distribution. *Theoretical Computer Science* 43, 169–188.
16. G.L. Jones, J.P. Hobert (2004): Sufficient burn-in for Gibbs samplers for a hierarchical random effects model. *Ann. Statist.* 32, pp. 784–817.
17. A.A. Johnson and G.L. Jones (2010): Gibbs sampling for a Bayesian hierarchical general linear model. *Electronic J. Statist.* 4, 313–333.
18. I. Kontoyiannis, L. Lastras-Montano, S.P. Meyn (2005): Relative Entropy and Exponential Deviation Bounds for General Markov Chains. *2005 IEEE International Symposium on Information Theory*.
19. K. Łatuszyński, B. Miasojedow and W. Niemiro (2009): Nonasymptotic bounds on the estimation error for regenerative MCMC algorithms. arXiv:0907.4915v1
20. K. Łatuszyński, W. Niemiro (2011): Rigorous confidence bounds for MCMC under a geometric drift condition. *J. of Complexity* 27, 23–38.
21. G. Lorden: On excess over the boundary. *Ann. Math. Statist.* 41, 520–527, 1970.
22. K.L. Mengersen, L.R. Tweedie (1996): Rates of convergence of the Hastings and Metropolis algorithms. *Ann. Statist.* 24, 1, 101–121.
23. S.P. Meyn and R.L. Tweedie: *Markov Chains and Stochastic Stability*. Springer-Verlag, 1993.
24. P. Mykland, L. Tierney and B. Yu (1995): Regeneration in Markov chain samplers. *J. Am. Statist. Assoc.*, 90, 233–241.
25. R. Neath, G.L. Jones (2009): Variable-at-a-time implementation of Markov chain Monte Carlo. Preprint. arXiv:0903.0664v1
26. W. Niemiro, P. Pokarowski (2009): Fixed precision MCMC Estimation by Median of Products of Averages. *J. Appl. Probab.* 46 (2), 309–329.

27. E. Nummelin (1978): A splitting technique for Harris recurrent Markov chains, *Z. Wahr. Verw. Geb.* 43, 309–318.
28. E. Nummelin (2002): MC's for MCMC'ists, *International Statistical Review*, 70, 215–240.
29. C.P. Robert and G. Casella: *Monte Carlo Statistical Methods*. Springer-Verlag, New York, 2004.
30. G.O. Roberts and J.S. Rosenthal (1997): Geometric ergodicity and hybrid Markov chains. *Elec. Comm. Probab.* 2 (2).
31. G.O. Roberts and J.S. Rosenthal (2004): General state space Markov chains and MCMC algorithms. *Probability Surveys* 1, 20–71.
32. J.S. Rosenthal (1995): Minorization conditions and convergence rates for Markov chains. *J. Amer. Statist. Association* 90, 558–566.
33. D. Rudolf (2008): Explicit error bounds for lazy reversible Markov chain Monte Carlo. *J. of Complexity*. 25, 11–24.
34. V. Roy, J.P. Hobert (2010): On Monte Carlo methods for Bayesian multivariate regression models with heavy-tailed errors. *J. Multivariate Anal.* 101, 1190–1202
35. D.B. Wilson (2000): How to couple from the past using a read-once source of randomness. *Random Structures Algorithms* 16 (1), 85–113.

Appendix

For convenience, we recall the two identities of Abraham Wald, which we need in the proof of Theorem 3.3. Proofs can be found e.g. in [6, Theorems 1 and 3 in Section 5.3].

Assume that $\eta_1, \dots, \eta_k, \dots$, are i.i.d. random variables and R is a stopping time such that $\mathbb{E}R < \infty$.

I Wald identity: If $\mathbb{E}|\eta_1| < \infty$ then

$$\mathbb{E} \sum_{k=1}^R \eta_k = \mathbb{E}R \mathbb{E} \eta_1.$$

II Wald identity: If $\mathbb{E} \eta_1 = 0$ and $\mathbb{E} \eta_1^2 < \infty$ then

$$\mathbb{E} \left(\sum_{k=1}^R \eta_k \right)^2 = \mathbb{E}R \mathbb{E} \eta_1^2.$$

In Section 5 we used the following formula taken from [28, Equation (4.1.4)]. In the notation of our Sections 2 and 3, for every $g \geq 0$ we have

$$\frac{\mathbb{E}_v \Xi_1(g)^2}{m} = \mathbb{E} \pi g(X_0)^2 + 2 \sum_{i=1}^{\infty} \mathbb{E} \pi g(X_0) g(X_i) \mathbb{I}(T > i). \quad (29)$$

In [28] this formula, with $g = \bar{f}$, is used to derive an expression for the asymptotic variance $\sigma_{\text{as}}^2(f) = \mathbb{E}_v \Xi_1(\bar{f})/m$ under the assumption that f is bounded. For $g \geq 0$, the proof is the same.